

# PRISMA: PROximal Iterative SMOothing Algorithm

Francesco Orabona

Toyota Technological Institute at Chicago  
francesco@orabona.com

Andreas Argyriou

Katholieke Universiteit Leuven  
andreas.argyriou@esat.kuleuven.be

Nathan Srebro

Toyota Technological Institute at Chicago  
nati@ttic.edu

March 8, 2013

## Abstract

Motivated by learning problems including max-norm regularized matrix completion and clustering, robust PCA and sparse inverse covariance selection, we propose a novel optimization algorithm for minimizing a convex objective which decomposes into three parts: a smooth part, a simple non-smooth Lipschitz part, and a simple non-smooth non-Lipschitz part. We use a time variant smoothing strategy that allows us to obtain a guarantee that does not depend on knowing in advance the total number of iterations nor a bound on the domain.

## 1 Introduction

We propose an optimization method (PRISMA — PROximal Iterative SMOothing Algorithm) for problems of the form<sup>1</sup>:

$$\min \{F(x) := f(x) + g(x) + h(x) : x \in \mathbb{R}^n\} \quad (1.1)$$

where:

- $f : \mathbb{R}^n \rightarrow \mathbb{R}$  is a convex and  $L_f$ -smooth function, that is, differentiable with a Lipschitz continuous gradient:  $\|\nabla f(x) - \nabla f(y)\| \leq L_f \|x - y\| \quad \forall x, y \in \mathbb{R}^n$ .

<sup>1</sup>All the theorems hold also in general Hilbert spaces, but for simplicity of exposition we consider a Euclidean setting.

- $g : \mathbb{R}^n \rightarrow \mathbb{R}$  is a convex  $\rho_g$ -Lipschitz continuous function:  $|f(x) - f(y)| \leq \rho_g \|x - y\|$ .
- $h : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$  is a proper, lower semicontinuous, convex (but possibly non-continuous/non-finite) function. For example,  $h$  could be an indicator function for a convex constraint.

We further assume that we can calculate gradients of  $f$ , and that  $g$  and  $h$  are “simple” in the sense that we can calculate their *proximity operators*<sup>2</sup>:

$$\text{prox}_g(x, \alpha) = \underset{u}{\operatorname{argmin}} \left\{ \frac{\|x - u\|^2}{2\alpha} + g(u) : u \in \mathbb{R}^n \right\},$$

where  $x \in \mathbb{R}^n$ ,  $\alpha \in \mathbb{R}_{++}$ , and similarly for  $h$ . That is, our method can be viewed as a black-box one, which accesses  $f, g$  and  $h$  only through the gradients of  $f$  and the proximity operators of  $g$  and  $h$ . Each iteration of PRISMA requires one evaluation of each of  $\nabla f$ ,  $\text{prox}_g$  and  $\text{prox}_h$  and a few vector operations of overall complexity  $O(n)$ , and after  $k$  such iterations we have

$$F(x_{k+1}) - F(x^*) = \mathcal{O} \left( \frac{L_f}{k^2} + \frac{\rho_g \log k}{k} \right).$$

**Applications.** Our main motivation for developing PRISMA was for solving optimization problems involving the matrix max-norm (aka  $\gamma_{2:\ell_1 \rightarrow \ell_\infty}$  norm). The max-norm has recently been shown to possess some advantages over the more commonly used trace-norm Lee et al. [2010], Jalali and Srebro [2012], Foygel and Srebro [2011], but is lagging behind it in terms of development of good optimization approaches—a gap we aim to narrow in this paper. Both norms are SDP representable, but standard SDP solvers are not applicable beyond very small scale problems. Instead, several first-order optimization approaches are available and commonly used for the trace-norm, including singular value thresholding Cai et al. [2008] and other Jaggi and Sulovsky [2010], Ma et al. [2009], Tomioka et al. [2010]. However, until now there have been no practical first-order optimization techniques available for the max-norm. In Sec. 4 we show how max-norm regularized problems can be written in the form (1.1) and solved using PRISMA, thus providing for the first time an efficient and practical first order method for max-norm optimization.

We also demonstrate how PRISMA can be applied also to other optimization problems with complex structure, including robust principal component analysis (robust PCA), sparse inverse covariance selection and basis pursuit. In particular, for robust PCA, we show how PRISMA yields better performance than previously published approaches, and for basis pursuit we obtain the best known convergence rate using only first-order and proximity oracles.

**Our Approach.** Following ideas of Nesterov Nesterov [2005b] and others, we propose to smooth the function  $g$  and use its proximity operator to obtain gradients of its smoothed version. However, unlike Nesterov [2005b], where a fixed amount of

<sup>2</sup>For more clarity here we explicitly define the proximity operator as a function of  $\alpha$  too because our algorithm heavily depends on this parameter changing over time.

smoothing is used throughout the run, we show how to gradually change the amount of smoothing at each iteration. We can then use an accelerated gradient descent approach on  $f$  plus the smoothed version of  $g$ , as in Nesterov [2005b]. We also use the ideas of partial linearization to handle the component  $h$ : instead of attempting to linearize it as in gradient descent, we include it as is in each iteration, as in FOBOS Duchi and Singer [2009] and ISTA/FISTA Beck and Teboulle [2009]. The gradual smoothing allows us to obtain a guarantee that does not depend on knowing in advance the total number of iterations, needed in Nesterov [2005b], nor a bound on the domain, needed in Nesterov [2005b], Chambolle and Pock [2011], and paying only an additional  $\log(k)$  factor.

**Notation.** We use  $f^*$  to denote the Fenchel conjugate,  $\delta_Q$  the indicator of the set  $Q$  (zero inside  $Q$  and infinite outside),  $\mathbf{S}_{++}^n$  the set of  $n \times n$  positive definite matrices, and  $\mathbf{S}_+^n$  the set of p.s.d. ones.

## 2 Smoothing of Nonsmooth Functions

As was suggested by Nesterov [2005a,b] and others, in order to handle the non-smooth component  $g$ , we approximate it using a smooth function. Such a smoothing plays a central role in our method, and we devote this section to carefully presenting it. In particular, we use the  $\beta$ -Moreau envelope Bauschke and Combettes [2011], known also as *Moreau-Yosida regularization*. For a function  $\varphi : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$ , the Moreau envelope  $\varphi_\beta$ , where  $\beta > 0$ , is defined as

$$\varphi_\beta(x) := \inf \left\{ \frac{\|x - u\|^2}{2\beta} + \varphi(u) : u \in \mathbb{R}^n \right\} \quad \forall x \in \mathbb{R}^n.$$

The function  $\varphi_\beta$  is a smooth approximation to  $\varphi$  (in fact, the best possible approximation of bounded smoothness), as summarized in the following Lemmas:

**Lemma 2.1** (Proposition 12.29 in Bauschke and Combettes [2011]). *Let  $\varphi : \mathbb{R}^n \rightarrow \mathbb{R}$  be a proper, lower semicontinuous, convex function and  $\beta > 0$ . Then  $\varphi_\beta$  is  $\frac{1}{\beta}$ -smooth and its gradient can be obtained from the proximity operator of  $\varphi$  as:  $\nabla(\varphi_\beta)(x) = \frac{1}{\beta}(x - \text{prox}_\varphi(x, \beta))$ .*

**Lemma 2.2.** *Let  $\varphi : \mathbb{R}^n \rightarrow \mathbb{R}$  a convex,  $\rho$ -Lipschitz function, then*

1. *if  $\beta > 0$ , then  $\varphi_\beta \leq \varphi \leq \varphi_\beta + \frac{1}{2}\beta\rho^2$ ;*
2. *if  $\beta \geq \beta' > 0$ , then  $\varphi_\beta \leq \varphi_{\beta'} \leq \varphi_\beta + \frac{1}{2}(\beta - \beta')\rho^2$ .*

**Proof.** The proof extends the one of Lemma 2.5 in Shalev-Shwartz et al. [2010] to vectorial functions. Define  $\Psi_x(u) = \frac{1}{2\beta}\|x - u\|^2 + \varphi(u)$ . We have  $\varphi_\beta(x) = \inf_u \Psi_x(u) \leq \Psi_x(x) = \varphi(x)$ , and this proves the left hand side of the first property. For the other side of the inequality, we have

$$\begin{aligned} \Psi_x(u) &= \frac{1}{2\beta}\|x - u\|^2 + \varphi(u) - \varphi(x) + \varphi(x) \\ &\geq \frac{1}{2\beta}\|x - u\|^2 + \varphi(x) - \rho\|x - u\| \end{aligned}$$

where we have used the Lipschitz property of  $\varphi$ . Hence

$$\begin{aligned}\varphi_\beta(x) &\geq \varphi(x) + \inf_u \left( \frac{\|x - u\|^2}{2\beta} - \rho|x - u| \right) \\ &= \varphi(x) - \frac{1}{2}\beta\rho^2.\end{aligned}$$

The second property follows from the first one and  $(\varphi_{\beta'})_{\beta-\beta'} = \varphi_\beta$  (Proposition 12.22 in Bauschke and Combettes [2011]).  $\blacksquare$

These two lemmas show how to obtain a smooth approximation of a Lipschitz continuous function, with the tradeoff between the smoothness and the degree of approximation controlled by the Lipschitz constant of the non-smooth function. Furthermore, in order to be able to work with the smoothed approximation  $\varphi_\beta$ , and in particular calculate its gradients, all that is required is access to the proximity operator of the non-smooth function  $\varphi$ . This is the reason we require access to  $\text{prox}_g$ .

**Relation to Nesterov Smoothing.** The Moreau envelope described above also underlies Nesterov's smoothing method Nesterov [2005b], although his derivation is somewhat different. Nesterov refers to a non-smooth function  $\varphi$  that can be written as:

$$\varphi(x) = \max\{\langle x, u \rangle - \hat{\varphi}(u) : u \in Q\} \quad (2.1)$$

where  $x \in \mathbb{R}^n$ ,  $Q \subseteq \mathbb{R}^m$  is a bounded closed convex set, and proposes to smoothen the function  $\varphi \circ A$ , where  $A : \mathbb{R}^n \rightarrow \mathbb{R}^m$  is a linear operator, with  $\tilde{\varphi}_\beta \circ A$ , where:

$$\tilde{\varphi}_\beta(x) := \max\{\langle x, u \rangle - \hat{\varphi}(u) - \beta d(u) : u \in Q\}, \quad (2.2)$$

with  $x \in \mathbb{R}^n$ ,  $\beta > 0$ , and  $d(u) \geq \frac{\sigma}{2}\|u - u_0\|^2$ . The smoothness and approximation properties of  $\tilde{\varphi}_\beta$  are then given in terms of the diameter of the set  $Q$ , the parameter  $\sigma$  and a certain norm of the matrix  $A$ . Any convex function  $\varphi$  can indeed be written in the form (2.1), where  $\hat{\varphi}$  is its convex conjugate and  $Q$  is the set of all subgradients of  $\varphi$  — the diameter of  $Q$  thus corresponds to the Lipschitz constant (bound on the magnitude of the subgradients) of  $\varphi$ . Focusing on  $A$  being the identity and  $d(u) = \frac{1}{2}\|u\|^2$  (i.e.  $\sigma = 1$ ), we have that  $\tilde{\varphi}_\beta$  defined in (2.2) is exactly the Moreau envelope  $\varphi_\beta$ :

**Proposition 2.3.** *If  $\varphi$  is a proper, lower semicontinuous convex function,  $\hat{\varphi}$  is its Fenchel conjugate,  $Q$  the set of all its subgradients and  $d(u) = \frac{1}{2}\|u\|^2$ , then  $\tilde{\varphi}_\beta = \varphi_\beta$ .*

This is obtained as a corollary from a slightly more general Lemma:

**Lemma 2.4.** *If  $\varphi$  is a proper, lower semicontinuous convex function which can be written as (2.1) and  $d(u) = \frac{1}{2}\|u\|^2$ , then  $\tilde{\varphi}_\beta = (\varphi \square \delta_Q^*)_\beta$ , where  $f \square g$  is the infimal convolution of  $f$  and  $g$ .*

**Proof.** From the definition of  $\varphi$  and [Bauschke and Combettes, 2011, Thm. 13.32], we derive that  $\hat{\varphi} = \varphi^* - \delta_Q$ . Hence

$$\begin{aligned}\tilde{\varphi}_\beta(x) &= \max_{u \in Q} \left\{ \langle x, u \rangle - \varphi^*(u) - \frac{\beta}{2}\|u\|^2 \right\} \\ &= \left( \varphi^* + \frac{\beta}{2}\|\cdot\|^2 + \delta_Q \right)^*(x) = (\varphi \square \delta_Q^*)_\beta(x),\end{aligned}$$

---

**Algorithm 1** PRISMA

---

**Parameters**  $\{\beta_k > 0 : k \in \mathbb{N}\}$ **Initialize**  $x_1 = y_1 \in \text{dom } h, \theta_1 = 1, L_1 = L_f + \frac{1}{\beta_1}$ **for**  $k = 1, 2, \dots$  **do**

$$L_{k+1} \leftarrow L_f + \frac{1}{\beta_{k+1}}$$

$$\theta_{k+1} \leftarrow \frac{2}{1 + \sqrt{1 + \frac{4L_{k+1}}{\theta_k^2 L_k}}}$$

$$x_{k+1} \leftarrow \text{prox}_h \left( \left(1 - \frac{1}{L_k \beta_k}\right) y_k - \frac{1}{L_k} \nabla f(y_k) + \frac{1}{L_k \beta_k} \text{prox}_g(y_k, \beta_k), \frac{1}{L_k} \right)$$

$$y_{k+1} \leftarrow x_{k+1} + \theta_{k+1} \left( \frac{1}{\theta_k} - 1 \right) (x_{k+1} - x_k)$$

**end for**

---

where the last equality follows from [Bauschke and Combettes, 2011, Thm. 13.32, Prop. 13.21 and Prop. 14.1].  $\blacksquare$

Nesterov's formulation is somewhat more general, allowing for different regularizers  $d(u)$  and different sets  $Q$ , but our presentation is a crisper and more direct statement of the assumptions on the function  $\varphi$  to be smoothed and its relationship to  $\varphi_\beta$ : we just need to be able to calculate the proximity operator of  $\varphi$ , to obtain gradients of  $\varphi_\beta$ , and we need to rely on its Lipschitz constant in order to be able to control the quality of the approximation.

### 3 PProximal Iterative SMOOTHing Algorithm

Our proposed method, PRISMA, is given as Algorithm 1. As is standard in “accelerated” first order methods (e.g. Beck and Teboulle [2009]), we keep track of two sequences of iterates,  $x_k$  and  $y_k$ . At each iteration we perform a proximal gradient update  $x_{k+1} \leftarrow \text{prox}_h \left( y_k - \frac{1}{L_k} \nabla (f + g_{\beta_k})(y_k), \frac{1}{L_k} \right)$ , where the gradient of  $g_{\beta_k}$  is calculated according to Lemma 2.1 and  $\beta_k$  is some sequence of parameters. We then set  $y_{k+1}$  to be a linear combination of  $x_{k+1}$  and  $x_k$ . Note also that the recursive formula for the sequence  $\theta_k$  satisfies

$$\frac{1}{\theta_{k+1}^2} - \frac{1}{\theta_{k+1}} = \frac{L_{k+1}}{L_k} \frac{1}{\theta_k^2}. \quad (3.1)$$

We establish the following guarantee on the values of the objective function.

**Theorem 3.1.** *If the sequence  $\beta_k$  is nonincreasing, then the iterates of Algorithm 1 satisfy,  $\forall x^* \in \text{dom } h$ ,*

$$F(x_{k+1}) - F(x^*) \leq \frac{1}{2} \theta_k^2 L_k \left( \|x^* - x_1\|^2 + \rho_g^2 \sum_{i=1}^k \frac{\beta_i}{\theta_i L_i} \right).$$

In the following we show two possible ways to set the sequences  $\beta_k$ . One possibility would be to decide on a fixed number of iterations  $T$  and to set  $\beta_k$  to a constant value.

**Corollary 3.2.** For a given  $T > 0$  set  $\beta_k = \frac{2\|x_1 - x^*\|}{\rho_g(T+1)} \forall k$ . Then for every  $x^* \in \text{dom } h$

$$F(x_{T+1}) - F(x^*) \leq 2 \frac{L_f \|x^* - x_1\|^2}{(T+1)^2} + 2 \frac{\rho_g \|x^* - x_1\|}{T+1}.$$

This choice is optimal if we know the number of iterations  $T$  we will use, as well as the Lipschitz constant and a bound on the distance from a minimizer, sharing the same limitation with Nesterov [2005b]. But with this fixed choice of smoothing, even if we perform additional iterations, we will not converge to the optimum of our original objective (only of the smoothed, thus approximate, objective).

We propose a dynamic setting of  $\beta_k$ , that gives a bound that holds uniformly for any number of iterations, and thus a method which converges to the optimum of the original objective.

**Corollary 3.3.** Let  $a > 0$  and set  $\beta_k = \frac{1}{ak}$  in Algorithm 1, then, uniformly for any  $k \in \mathbb{N}$  and for every  $x^* \in \text{dom } h$

$$F(x_{k+1}) - F(x^*) \leq 2 \frac{L_f + ak}{(k+1)^2} \left[ \|x^* - x_1\|^2 + \frac{\rho_g^2}{a} \left( \frac{3}{2a} \log \frac{L_f + ak}{L_f + a} + \frac{1}{L_f + a} \right) \right].$$

**Proof.** It is easy to derive that  $\frac{1}{k+1} < \theta_k \leq \frac{2}{k+1}$  (by induction) and the elementary chain of inequalities  $\sum_{i=1}^k \frac{1}{a+bi} = \frac{1}{a+b} + \sum_{i=2}^k \frac{1}{a+bi} \leq \frac{1}{a+b} + \int_1^k \frac{1}{a+bx} dx$ . The assertion then follows from Theorem 3.1 and these facts. ■

Note that the optimal choice of  $a$  does depend on the Lipschitz constant  $\rho_g$  and on the distance of our initial point from a minimizer, but that the algorithm does converge for *any* choice of  $a$ . Thus, compared to the bound in Corollary 3.2, the price we pay for not knowing in advance how many iterations will be needed is only an additive logarithmic factor<sup>3</sup>.

When  $f = 0$  the updates reduce to  $x_{k+1} \leftarrow \text{prox}_h(\text{prox}_g(y_k, \frac{1}{ak}), \frac{1}{ak})$ , and  $\theta_k = \frac{1}{k}$ . This can be viewed as an adaptive version of an accelerated backward-backward splitting algorithm.

Before proving Theorem 3.1, we need two additional technical lemmas.

**Lemma 3.4** (*descent lemma*, Thm. 18.15 in Bauschke and Combettes [2011], Thm. 2.1.5 in Nesterov [2004]). *The convex function  $f$  is  $L$ -smooth if and only if*

$$f(x) \leq f(y) + \langle x - y, \nabla f(y) \rangle + \frac{L}{2} \|x - y\|^2 \quad \forall x, y \in \mathbb{R}^n.$$

<sup>3</sup>Ouyang and Gray [2012a] use a similar proof technique, for the simpler case when  $h = 0$ , and they report a convergence rate bound without the logarithmic term. Unfortunately their proof contains an error; they fixed it in the arxiv version of their paper Ouyang and Gray [2012b]. Their new theorem contains the term  $\max_k \|x^* - x_k\|^2$ , which is not known how it can be bounded.

**Lemma 3.5** (Thm. 2.1.2 in Nesterov [2004]). *For a  $\mu$ -strongly convex function  $g$ , and  $\hat{v} = \operatorname{argmin} g$ ,*

$$g(w) - g(\hat{v}) \geq \frac{\mu}{2} \|w - \hat{v}\|^2 \quad \forall w \in \mathbb{R}^n.$$

**Proof. (of Theorem 3.1)** Denote  $F^{\beta_k}(x) := f(x) + g_{\beta_k}(x) + h(x)$ . Fix an arbitrary  $k \in \mathbb{N}$ . Since  $f + g_{\beta_k}$  is  $L_k$ -smooth by Lemma 2.1, applying Lemma 3.4 yields

$$\begin{aligned} F^{\beta_k}(x_{k+1}) &\leq f(y_k) + g_{\beta_k}(y_k) + \frac{L_k}{2} \|x_{k+1} - y_k\|^2 \\ &\quad + \langle \nabla(f + g_{\beta_k})(y_k), x_{k+1} - y_k \rangle + h(x_{k+1}). \end{aligned}$$

We now use the strong convexity of the function  $x \mapsto \langle \nabla(f + g_{\beta_k})(y_k), x \rangle + \frac{L_k}{2} \|x - y_k\|^2 + h(x)$ , whose minimizer equals  $x_{k+1}$  (by Lemma 2.1). Using Lemma 3.5 with  $w = (1 - \theta_k)x_k + \theta_k x$ , we obtain

$$\begin{aligned} F^{\beta_k}(x_{k+1}) &\leq f(y_k) + g_{\beta_k}(y_k) + h((1 - \theta_k)x_k + \theta_k x) \\ &\quad + \langle \nabla(f + g_{\beta_k})(y_k), (1 - \theta_k)x_k + \theta_k x - y_k \rangle \\ &\quad + \frac{L_k}{2} \|(1 - \theta_k)x_k + \theta_k x - y_k\|^2 - \frac{L_k}{2} \|(1 - \theta_k)x_k + \theta_k x - x_{k+1}\|^2. \end{aligned}$$

We let  $z_k := \frac{1}{\theta_k} y_k + \left(1 - \frac{1}{\theta_k}\right) x_k$  for every  $k \in \mathbb{N}$  and note that  $z_{k+1} = x_k + \frac{1}{\theta_k}(x_{k+1} - x_k)$  by the algorithmic construction of  $y_{k+1}$ . Therefore,

$$\begin{aligned} F^{\beta_k}(x_{k+1}) &\leq f(y_k) + g_{\beta_k}(y_k) + h((1 - \theta_k)x_k + \theta_k x) \\ &\quad + \langle \nabla(f + g_{\beta_k})(y_k), (1 - \theta_k)x_k + \theta_k x - y_k \rangle \\ &\quad + \frac{\theta_k^2 L_k}{2} \|x - z_k\|^2 - \frac{\theta_k^2 L_k}{2} \|x - z_{k+1}\|^2 \\ &\leq f(y_k) + g_{\beta_k}(y_k) + (1 - \theta_k)h(x_k) + \theta_k h(x) \\ &\quad + \langle \nabla(f + g_{\beta_k})(y_k), (1 - \theta_k)x_k + \theta_k x - y_k \rangle \\ &\quad + \frac{\theta_k^2 L_k}{2} \|x - z_k\|^2 - \frac{\theta_k^2 L_k}{2} \|x - z_{k+1}\|^2 \\ &\leq (1 - \theta_k)F^{\beta_k}(x_k) + \theta_k F^{\beta_k}(x) \\ &\quad + \frac{\theta_k^2 L_k}{2} \|x - z_k\|^2 - \frac{\theta_k^2 L_k}{2} \|x - z_{k+1}\|^2, \end{aligned}$$

where we have used the convexity of the functions  $f, g, h$ . Using Property 1 in Lemma 2.2 we have

$$\begin{aligned} F^{\beta_k}(x_{k+1}) - F(x) &\leq \frac{\theta_k^2 L_k}{2} (\|x - z_k\|^2 - \|x - z_{k+1}\|^2) \\ &\quad + (1 - \theta_k) (F^{\beta_k}(x_k) - F(x)). \end{aligned}$$

We now use Property 2 in Lemma 2.2 to change the smoothing parameter. Denoting by  $D_k = F^{\beta_k}(x_k) - F(x)$ , we obtain

$$\frac{D_{k+1}}{\theta_k^2 L_k} \leq \frac{1}{2} (\|x - z_k\|^2 - \|x - z_{k+1}\|^2) + \frac{(\beta_k - \beta_{k+1})\rho_g^2}{2\theta_k^2 L_k} + \frac{1 - \theta_k}{\theta_k^2 L_k} D_k.$$

Using the definition of  $\theta_k$  (3.1), and summing we obtain

$$\frac{D_{k+1}}{\theta_k^2 L_k} \leq \frac{1}{2} (\|x - z_1\|^2 - \|x - z_{k+1}\|^2) + \sum_{i=1}^k \frac{(\beta_i - \beta_{i+1})\rho_g^2}{2\theta_i^2 L_i}.$$

Finally, we apply Property 1 in Lemma 2.2 to obtain  $F(x_{k+1}) \leq F^{\beta_{k+1}}(x_{k+1}) + \frac{1}{2}\beta_{k+1}\rho_g^2$  and reordering the terms we obtain that

$$D_{k+1} \leq \frac{1}{2}\theta_k^2 L_k \|x - x_1\|^2 + \frac{1}{2}\beta_{k+1}\rho_g^2 + \frac{1}{2}\theta_k^2 L_k \sum_{i=1}^k \frac{(\beta_i - \beta_{i+1})\rho_g^2}{\theta_i^2 L_i}.$$

Applying the definition of  $\theta_k$  (3.1) and gathering the  $\beta_i$  terms, we obtain the stated bound.  $\blacksquare$

## Relation to Prior Work

With  $g = 0$ , PRISMA with fixed smoothing reduces to a variant of FISTA [Beck and Teboulle, 2009], with essentially the same guarantee of  $\mathcal{O}(\frac{L_f}{k^2})$ .

Conversely, with  $h = 0$ , or  $h$  being an indicator for a convex domain, PRISMA becomes similar to Nesterov's smoothed accelerated method [Nesterov, 2005b], obtaining the rate of  $\mathcal{O}(\frac{\log k}{k})$ , but with some important differences. First, we provide an explicit bound in terms of the Lipschitz constant of  $g$ , as discussed in Sec. 2. Second, Nesterov's method relies on a fixed domain and involves two projections onto the domain at each iteration, whereas PRISMA uses only a single projection. Moreover, Nesterov's methods and guarantees rely on the domain being bounded and require that the number of iterations must be fixed in advance, in order to set its parameters. On the other hand, the dynamic tuning of  $\beta_k$  in PRISMA allows us to obtain a guarantee that depends only on  $\|x^* - x_1\|$ , and with parameters (i.e.  $\beta_k$  and  $\theta_k$ ) which do not have to depend on it, or on a fixed number of iterations. Unbounded domains are frequently encountered in practice, for example, in our motivating application of max-norm optimization. The excessive gap primal-dual algorithm in Nesterov [2005a] improves over the one in Nesterov [2005b] because it does not need to fix in advance the number of iterations, but it shares the same shortcoming on the assumption of the bounded domain.

The primal-dual algorithm in Chambolle and Pock [2011] can minimize functions of the form  $g(Kx) + h(x)$ , where  $K$  is a linear operator, and it assumes to have access to  $\text{prox}_g$  and  $\text{prox}_h$ . The original problem is transformed into a saddle point problem and they obtain a  $\mathcal{O}(\frac{1}{k})$  convergence rate for the primal-dual gap, for non-smooth objective functions. However this kind of guarantee will translate to a guarantee on the functional objective value only assuming that  $g$  and  $h^*$  are Lipschitz (see also the discussion on



this point in Loris and Verhoeven [2011]). Moreover it is not clear how the obtained bound explicitly depends on the Lipschitz constants and the other relevant quantities.

A different approach is given by the Alternating Direction Method of Multipliers (ADMM) algorithms [Boyd et al., 2011], that optimize the augmented Lagrangian associated to the optimization problem, updating alternatively the variables of the problem. ADMM approaches have the drawback that constraints in the optimization problem are not necessarily satisfied at every iteration. This makes them ill suited to problems like max-norm matrix completion, where two matrices would be generated at each iteration, only one of the two being PSD, and for none of the two a convergence rate bound is known. Recently, Goldfarb et al. [2012], Scheinberg et al. [2010] presented an ADMM method, called Alternating Linearization Method (ALM), which uses alternate projections of an augmented Lagrangian. It is applicable when  $g = 0$ , that is, for objectives which can be decomposed into a smooth plus a “simple” function, and it has the same rate of convergence of FISTA. It can also be used with non-smooth functions, using the smoothing method proposed by Nesterov [2005b], and hence it shares the same shortcomings.

Another recent related method [Loris and Verhoeven, 2011] is applicable when  $h = 0$  and  $f$  is quadratic. This method differs significantly from both PRISMA and Nesterov’s aforementioned methods, in part because no  $\theta_k$  and no affine update of  $y_k$  are required. The applicability of Loris and Verhoeven [2011] is not as general as PRISMA (for example it does not apply to max norm regularization or basis pursuit) and it is not clear whether it can be extended to any smooth function  $f$ . Similar to PRISMA, it attains an  $\mathcal{O}(\frac{1}{k})$  rate without requiring boundedness of the domain.

PRISMA combines the benefits of many of the above methods: we can handle functions decomposable to three parts, as in the max-norm optimization problem, with an unbounded domain, and without having the parameters to depend on the distance to a minimizer nor to fix a priori the number of iterations, at the cost of only an additional log-factor. This is an important advantage in practice: setting parameters is tricky and we typically do not know in advance how many iterations will be required.

## 4 Applications

### 4.1 Max Norm Regularized Matrix completion

The max norm (also known as the  $\gamma_{2:\ell_1 \rightarrow \ell_\infty}$  norm) of a matrix  $X$  is given by  $\|X\|_{max} = \inf_{X=UV^\top} \|U\|_{2,\infty} \|V\|_{2,\infty}$ , where  $\|U\|_{2,\infty}$  is the maximal row norm of  $U$ . The max norm has been suggested as an effective regularizer for matrix completion (i.e. collaborative filtering) problems, with some nice theoretical and empirical advantages over the more commonly used nuclear norm (or trace norm) Srebro et al. [2005], Foygel and Srebro [2011], Lee et al. [2010]. It has also been recently suggested to use it in clustering problems, where it was shown to have benefits over a nuclear-norm based approach, as well as over spectral clustering Jalali and Srebro [2012]. However, the max norm remains much less commonly used than the nuclear norm. One reason might be the lack of good optimization approaches for the max norm. It has also been recently brought our attention that Jaggi [2011] suggested an optimization approach

for  $\min_{\|X\|_{max} \leq 1} f(X)$  which is similar in some ways to these first order methods, and which requires  $\mathcal{O}\left(\frac{L_f}{k}\right)$  iterations, but where each iteration consists of solving a max-cut type problem instead of an SVD—the practicality of such an approach is not yet clear and as far as we know it has not been implemented. Instead, the only practical approach we are aware of is a non-convex proximal point (NCP) approach, without any performance guarantees, and relying on significant “tweaking” of learning rates Lee et al. [2010]. Here we present the first practical first order optimization method for max norm regularized problems.

To demonstrate the method, we focus on the matrix completion setting, where we are given a subset  $\Omega$  of observed entries of a matrix  $M \in \mathbb{R}^{m \times n}$  and would like to infer the unobserved entries. Using max norm regularization, we have to solve the following optimization problem:

$$\min_W \lambda \|W\|_{max} + \|P_\Omega(W) - P_\Omega(M)\|_F^2. \quad (4.1)$$

where  $P_\Omega$  is a projection onto the entries in  $\Omega$  (zeroing out all other entries). Expressing the max norm through its semi-definite representation we can rewrite (4.1) as

$$\begin{aligned} \min_{A, B, X} \quad & \left\{ \lambda \max \text{diag} \left( \begin{smallmatrix} A & X \\ X^\top & B \end{smallmatrix} \right) + \|P_\Omega(X) - P_\Omega(M)\|_F^2 \right. \\ & \left. + \delta_{\mathbf{S}_+^{m+n}} \left( \begin{smallmatrix} A & X \\ X^\top & B \end{smallmatrix} \right) \right\}. \end{aligned}$$

Setting  $f(X) = \|P_\Omega(X) - P_\Omega(M)\|_F^2$ ,  $g(X) = \lambda \max \text{diag}(X)$  and  $h(X) = \delta_{\mathbf{S}_+^{m+n}}(X)$ , we can apply PRISMA. The proximity operator of  $h$  corresponds to the projection onto the set of  $(m+n) \times (m+n)$  positive semidefinite matrices, that can be computed by an eigenvalue decomposition and setting to zero all the negative eigenvalues. The proximity operator of  $g$  amounts to thresholding the diagonal elements of the matrix.

To demonstrate the applicability of PRISMA, we conducted runtime experiments on the MovieLens 100K Dataset<sup>4</sup>, selecting an equal and increasing number of users and movies with the most ratings. We used training data equal to 25% of the total number of entries, selected uniformly at random among the ranked entries. We implemented PRISMA in MATLAB. We stopped PRISMA when the relative difference  $\frac{\|X_{k+1} - X_k\|_F}{\|X_k\|_F}$  was less than  $10^{-5}$ . The computation bottleneck of the algorithm is the computation of the eigenvalue decomposition at each step. So, to speed-up the algorithm, we implemented the strategy to calculate only the top  $c$  eigenvalues, where  $c$  is defined as  $\min(m+n, (\# \text{ eigenvalues} > 0 \text{ at previous iteration}) + 1)$ . If the smallest calculated eigenvalue is positive, we increase  $c$  by 5 and recalculate them, otherwise we proceed with the projection step. This strategy is based on the empirical observation that the algorithm produces solutions of decreasing rank, and practically this procedure gives a big speed-up to the algorithm. This kind of strategy is common in the matrix completion literature and RPCA, see e.g. Goldfarb et al. [2009]. We compared our algorithm to SDPT3 Tütüncü et al. [2003], a state-of-the package for semidefinite-quadratic-linear programming. For simplicity the parameter  $\lambda$  was set the same in all

<sup>4</sup><http://www.grouplens.org/node/73>

Matrix Size	SDPT3		PRISMA		
	Cost	Time	Cost	Iter.	Time
100x100	1078.76	13s	1079.00	7366	124s
150x150	2617.65	89s	2618.11	7866	218s
200x200	4795.76	364s	4796.97	8908	281s
250x250	7736.76	1323s	7738.59	9524	404s
300x300	—	—	11406.39	9843	578s

Table 1: Performance evaluation of PRISMA and SDPT3 on subsets of the MovieLens 100K dataset.

the experiments, equal to  $0.2|\Omega|$  that guarantees reasonable good generalization performance for all the sizes considered. The parameter  $a$  of PRISMA was set to a rough estimate of the optimal value, that is, to  $\frac{\lambda\sqrt{|\Omega|}}{(m+n)\|P_\Omega(M)\|_F}$ . The running times and the final value of the objective function are listed in Table 1. SDPT3, a highly optimized SDP solver relying on interior point methods, indeed has very good performance for relatively small problems. But very quickly, the poor scaling of interior point methods kick in, and the runtime increases dramatically. More importantly, SDPT3 ran out of memory and could not run on problems of size 300x300 or larger, as did other SDP solvers we tried. On the other hand, the number of iterations required by PRISMA to reach the required precision increases only mildly with the size of the matrices and the increase in time is mainly due to the more expensive SVDs. We could run PRISMA on much larger problems, including the full MovieLens 100K data set. In Figure 1 we report the objective value as function of time (in seconds) for PRISMA on this data set. The time of each PRISMA iteration decreased over time, as the rank of the iterate decreases, ranging from 30 seconds of the first iterations to 2 seconds of the last ones, yielding a total runtime of several hours on the entire dataset. In contrast, it is inconceivable to use generic SDP methods for such data. On the other extreme, non-convex proximal-point (NCPP) methods are indeed much faster on this (and even much larger) data sets, taking only about ten seconds to reach a reasonable solution. However, even running NCPP for many hours with different step size settings and regardless of the allowed dimensionality, NCPP was not able to reach the global minimum, and was always at least 4% worst than the solution found by PRISMA—the minimum objective value attained by NCPP is also plotted in Figure 1. PRISMA thus fills a gap between the very pricey SDP methods that are practical only on tiny data sets, and the large-scale NCPP methods which require significant parameter tweaking and produce only rough (even if sometimes satisfying) solutions.

## 4.2 Robust PCA

In the method of Candes et al. [2011] for robust PCA, a data matrix  $M \in \mathbb{R}^{n_1 \times n_2}$  is given and the goal is to solve

$$\min \{ \|W\|_{tr} + \lambda \|S\|_1 : W + S = M, W, S \in \mathbb{R}^{n_1 \times n_2} \}.$$

This problem is of the form (1.1) with the choice  $f = 0, g(W) = \lambda \|M - W\|_1, h(W) = \|W\|_{tr}$ . We used the RPCA method to solve the task of background extraction from

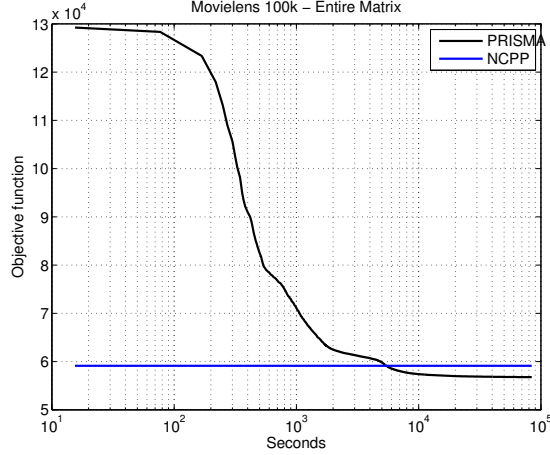


Figure 1: Objective value for PRISMA as a function of iterations in log scales.

surveillance video. By stacking the columns of each frame into a long vector, we get a matrix  $M$  whose columns correspond to the sequence of frames of the video. This matrix  $M$  can be decomposed into the sum of two matrices  $M = W + S$ . The matrix  $W$ , which represents the background in the frames, should be of low rank due to the correlation between frames. The matrix  $S$ , which represents the moving objects in the foreground in the frames, should be sparse since these objects usually occupy a small portion of each frame.

We have compared our MATLAB implementation of PRISMA to ALM Goldfarb et al. [2009, 2012], where this problem is solved with an augmented Lagrangian method, smoothing both the trace and the  $\ell_1$  norm. We further used the ALM continuation strategy described in Goldfarb et al. [2009], that is usually beneficial in practice even though no theoretical iteration count guarantees are known for it. The parameter  $a$  of PRISMA, analogously to max-norm experiments, is set to  $\frac{\lambda\sqrt{n_1 n_2}}{\|M\|_F}$ . We also compared it to the static strategy, setting  $\beta_k$  to constant times  $a$ .

In our experiments, we used two videos<sup>5</sup> introduced in Li et al. [2004], “Hall” and “Campus” with the same parameters used in Goldfarb et al. [2009]. Since for all algorithms the complexity for each update is roughly the same, and dominated by an SVD computation, we directly compare the number of iterations in order to reduce the hardware/implementation dependence of the experiments. The average CPU time per iteration for the algorithms and datasets are reported in Table 2. The results are shown in Figure 2. On both image sequences PRISMA outperforms the static algorithm, uniformly over the number of iterations. Notice also that for some constant settings the algorithm does not converge to the optimum, as expected from Theorem 3.1. Moreover it also outperforms the state-of-the-art algorithm ALM.

<sup>5</sup>[http://perception.i2r.a-star.edu.sg/bk\\_model/bk\\_index.html](http://perception.i2r.a-star.edu.sg/bk_model/bk_index.html)

	Campus	Hall
PRISMA (static & dynamic)	$3.5 \pm 0.1$	$1.9 \pm 0.01$
ALM	$5 \pm 0.3$	$2.66 \pm 0.1$

Table 2: Time (in seconds) per step, average  $\pm$  standard deviation.

### 4.3 Sparse Inverse Covariance Selection

In sparse inverse covariance selection (SICS) – see Wainwright et al. [2007], Yuan and Lin [2007] and references therein – the goal is to solve

$$\min_{X \in \mathbf{S}_{++}^n} -\log \det(X) + \langle \Sigma, X \rangle + \lambda \|X\|_1, \quad (4.2)$$

where  $\Sigma \in \mathbf{S}_{++}^n$  is a given positive semidefinite matrix. This problem is of the form (1.1) with the choice  $f = \langle \Sigma, \cdot \rangle$ ,  $g = \lambda \|\cdot\|_1$ ,  $h = -\log \det(\cdot) + \delta_{\mathbf{S}_{++}^n}(\cdot)$ . The required proximity operator of  $h$  at  $X \in \mathbf{S}_{++}^n$  can be easily computed by computing a singular value decomposition of  $X$  and solving a simple quadratic equation to obtain each singular value of the result. We have compared PRISMA with two state-of-the-art approaches: ALM Scheinberg et al. [2010], which achieves state-of-the-art empirical results, and ADMM Boyd et al. [2011]<sup>6</sup>, a recently popular general optimization method that performs very well in practice, but lacks iteration complexity guarantees.

ALM uses a complex continuation strategy that requires tuning of four different parameters – we used the data-set specific parameter settings provided by Scheinberg et al. [2010]. For ADMM we used the default parameters. For both ALM and ADMM there is no clear theory on how to set their parameters. To set the  $a$  parameter of PRISMA we need a rough estimate of  $\|X^*\|_F$ . The optimal solution of (4.2) will be very close to a diagonal matrix, so we approximate its optimum of adding the constraint  $X = \alpha I$ , and it is easy to see that the optimal  $\alpha$  is  $\frac{1}{1+\lambda}$ . Using the fact that  $\rho_g = \lambda n$ , we set  $a = \lambda(1 + \lambda)\sqrt{n}$  in all experiments. That is, both for ADMM and PRISMA we did not specifically tune any parameters.

We experimented on same five gene expression network data sets from Li and Toh [2010], which were also used by Scheinberg et al. [2010], and used the same setting  $\lambda = 0.5$ .

The results on the datasets are shown in Fig. 3-7. Again we show the objective function in relation to the number of iterations, being the complexity per step of the algorithms roughly the same. The average times per iteration are reported in Table 3,4. We can see that the performance of PRISMA is pretty close to the one of ALM, while ADMM results in a slower convergence. Notice that the difference between ALM and ADMM is probably mainly due to the continuation strategy used in ALM, whose parameters are been tuned on these datasets. On the other hand PRISMA just has one parameter, whose optimal setting is given by the theory.

<sup>6</sup>We used the MATLAB code available at <http://www.stanford.edu/~boyd/papers/admm/>.

	Hereditary breast cancer	Lymph
PRISMA	$11.74 \pm 0.75$	$0.33 \pm 0.02$
ALM	$7.5 \pm 0.09$	$0.27 \pm 0.05$
ADMM	$13 \pm 0.23$	$0.47 \pm 0.06$

Table 3: Time (in seconds) per step, average  $\pm$  standard deviation.

	Estrogen receptor	Arabidopsis thaliana	Leukemia
PRISMA	$0.72 \pm 0.04$	$1.15 \pm 0.09$	$3.61 \pm 0.19$
ALM	$0.45 \pm 0.03$	$0.75 \pm 0.04$	$2.46 \pm 0.05$
ADMM	$0.80 \pm 0.02$	$1.40 \pm 0.05$	$4.15 \pm 0.09$

Table 4: Time (in seconds) per step, average  $\pm$  standard deviation.

## 5 Discussion and Future work

We have proposed PRISMA, a new algorithm which can be used to solve many convex nonsmooth optimization problems in machine learning, signal processing and other areas. PRISMA belongs to the broad family of proximal splitting methods and extends in different ways forward-backward splitting, accelerated proximal, smoothing and excessive gap methods. PRISMA is distinguished from these methods by its much wider applicability, which allows for solving problems such as basis pursuit or semidefinite programs such as max norm matrix completion.

We have validated our method with experiments on matrix completion, robust PCA problems, and sparse inverse covariance selection, showing that a simple to code implementation of PRISMA can handle large scale problems and outperforms or be equal in efficiency to current state of the art optimization algorithms. But PRISMA can be useful also for other learning-related problems.

For example, Basis Pursuit (BP) Chen et al. [2001], Guigue et al. [2005] can also be solved with PRISMA. The BP optimization problem is formulated as

$$\min \{ \|x\|_1 : Ax = b, x \in \mathbb{R}^d \},$$

where  $A \in \mathbb{R}^{m \times d}$ ,  $b \in \mathbb{R}^m$  are prescribed input/output data. In most applications, the sample size is much smaller than the dimensionality,  $m \ll d$ . This problem is of the form (1.1) with the choice  $f = 0$ ,  $g = \|\cdot\|_1$ ,  $h = \delta_{\mathcal{A}}$ , where  $\mathcal{A} = \{x \in \mathbb{R}^d : Ax = b\}$ . BP can be rephrased as a linear program, which is then solved with standard approaches (simplex, interior point methods, etc.), or it can be solved with ADMM, but in both cases no complexity results are known. It is easy to see that, for every  $x \in \mathbb{R}^d$ , its projection on  $\mathcal{A}$  can be written as  $\text{Proj}_{\mathcal{A}}(x) = x - A^\top z$ , where  $z$  is any solution of  $AA^\top z = Ax - b$ , hence we can use PRISMA to solve this problem. The complexity of the projection step is  $\mathcal{O}(dm)$ , since  $m \ll d$ . It would be interesting to see whether PRISMA would indeed yield empirical advantages here, as well as in varied other learning and optimization problems.

## References

- H. H. Bauschke and P. L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. CMS Books in Mathematics. Springer, 2011.
- A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal of Imaging Sciences*, 2(1):183–202, 2009.
- S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.
- J. F. Cai, E. J. Candes, and Z. Shen. A singular value thresholding algorithm for matrix completion. *SIAM Journal on Optimization*, 20(4):1956–1982, 2008.
- E. J. Candes, X Li, Y. Ma, and J. Wright. Robust principal component analysis? *Journal of the Association for Computing Machinery*, 58(3), 2011.
- A. Chambolle and T. Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145, 2011.
- S. S. Chen, D. L. Donoho, and M. A. Saunders. Atomic decomposition by basis pursuit. *SIAM Review*, 43(1):129–159, 2001.
- J. Duchi and Y. Singer. Efficient online and batch learning using forward backward splitting. *The Journal of Machine Learning Research*, 10:2899–2934, 2009.
- R. Foygel and N. Srebro. Concentration-based guarantees for low-rank matrix reconstruction. *24th Annual Conference on Learning Theory (COLT)*, 2011.
- D. Goldfarb, S. Ma, and K. Scheinberg. Fast alternating linearization methods for minimizing the sum of two convex functions. arXiv:0912.4571, 2009.
- D. Goldfarb, S. Ma, and K. Scheinberg. Fast alternating linearization methods for minimizing the sum of two convex functions. *Mathematical Programming*, pages 1–34, 2012.
- V. Guigue, A. Rakotomamonjy, and S. Canu. Kernel basis pursuit. In *Proceedings of the Sixteenth European Conference on Machine Learning*, pages 146–157. Springer, 2005.
- M. Jaggi. Convex optimization without projection steps. arXiv:1108.1170, 2011.
- M. Jaggi and M. Sulovsky. A simple algorithm for nuclear norm regularized problems. In *Proceedings of the 27th International Conference on Machine Learning (ICML 2010)*, pages 471–478, 2010.
- A. Jalali and N. Srebro. Clustering using max-norm constrained optimization. In *International Conference on Machine Learning*, 2012.

- J. Lee, B. Recht, R. Salakhutdinov, N. Srebro, and J. Tropp. Practical large-scale optimization for max-norm regularization. In *Advances in Neural Information Processing Systems 23*, pages 1297–1305, 2010.
- L. Li and K.-C. Toh. An inexact interior point method for  $\ell_1$ -regularized sparse covariance selection. 2010.
- L. Li, W. Huang, I. Y. Gu, and Q. Tian. Statistical modeling of complex backgrounds for foreground object detection. *IEEE Trans Image Process*, 13(11):1459–1472, 2004.
- I. Loris and C. Verhoeven. On a generalization of the iterative soft-thresholding algorithm for the case of non-separable penalty. *Inverse Problems*, 27(12):125007, 2011.
- S. Ma, D. Goldfarb, and L. Chen. Fixed point and Bregman iterative methods for matrix rank minimization. *Mathematical Programming*, pages 1–33, 2009.
- Y. Nesterov. *Introductory lectures on convex optimization: A basic course*. Springer, 2004.
- Y. Nesterov. Excessive gap technique in nonsmooth convex minimization. *SIAM J. on Optimization*, 16(1):235–249, May 2005a.
- Y. Nesterov. Smooth minimization of non-smooth functions. *Mathematical Programming*, 103(1):127–152, 2005b.
- H. Ouyang and A. Gray. Stochastic smoothing for nonsmooth minimizations: Accelerating SGD by exploiting structure. In *Proceedings of the 29th International Conference on Machine Learning (ICML 2012)*, 2012a.
- H. Ouyang and A. Gray. Stochastic smoothing for nonsmooth minimizations: Accelerating SGD by exploiting structure. arXiv:1205.4481, 2012b.
- K. Scheinberg, S. Ma, and D. Goldfarb. Sparse inverse covariance selection via alternating linearization methods. In *Advances in Neural Information Processing Systems 23*, pages 2101–2109. 2010.
- S. Shalev-Shwartz, N. Srebro, and T. Zhang. Trading accuracy for sparsity in optimization problems with sparsity constraints. *SIAM Journal on Optimization*, 20(6), 2010.
- N. Srebro, J. D. M. Rennie, and T. S. Jaakkola. Maximum-margin matrix factorization. In *Advances in Neural Information Processing Systems 17*, pages 1329–1336. MIT Press, 2005.
- R. Tomioka, T. Suzuki, M. Sugiyama, and H. Kashima. A fast augmented Lagrangian algorithm for learning low-rank matrices. In *International Conference on Machine Learning*, 2010.
- Reha H. Tütüncü, K. C. Toh, and Michael J. Todd. Solving semidefinite-quadratic-linear programs using sdpt3. *Math. Program.*, 95(2):189–217, 2003.



- M. J. Wainwright, P. Ravikumar, and J. D. Lafferty. High-dimensional graphical model selection using  $\ell_1$ -regularized logistic regression. In *Advances in Neural Information Processing Systems 19*, pages 1465–1472, 2007.
- M. Yuan and Y. Lin. Model selection and estimation in the Gaussian graphical model. *Biometrika*, 94(1):19–35, 2007.

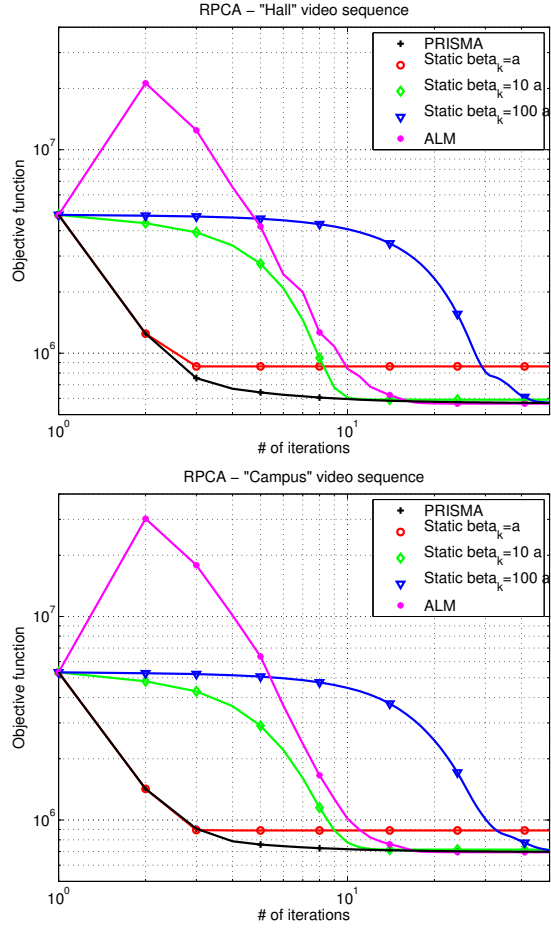


Figure 2: Comparison of the objective values as a function of the number of iterations, in log scales.

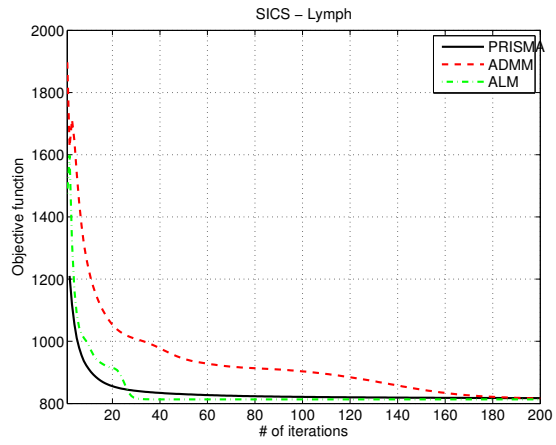


Figure 3: Comparison of the objective values as a function of the iterations.

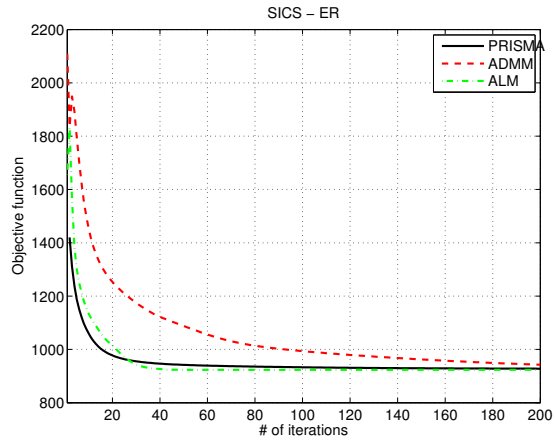


Figure 4: Comparison of the objective values as a function of the iterations.

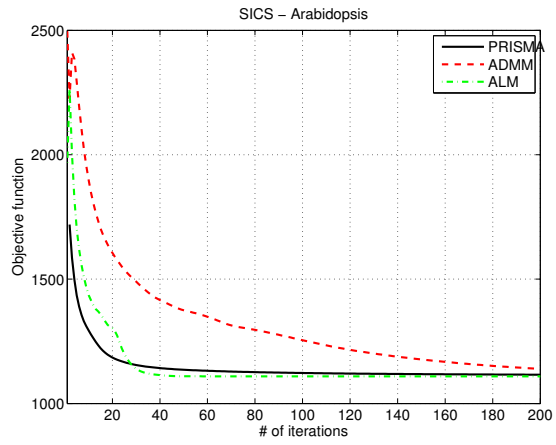


Figure 5: Comparison of the objective values as a function of the iterations.

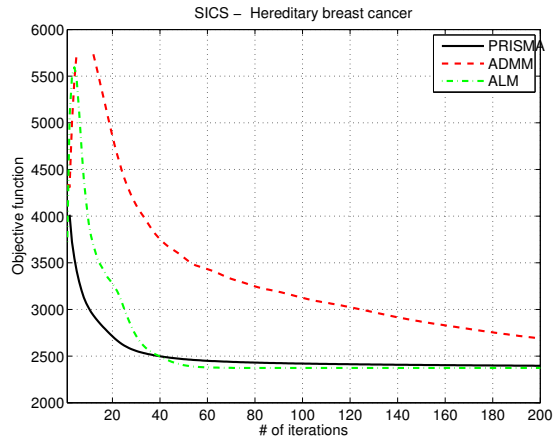


Figure 6: Comparison of the objective values as a function of the iterations.

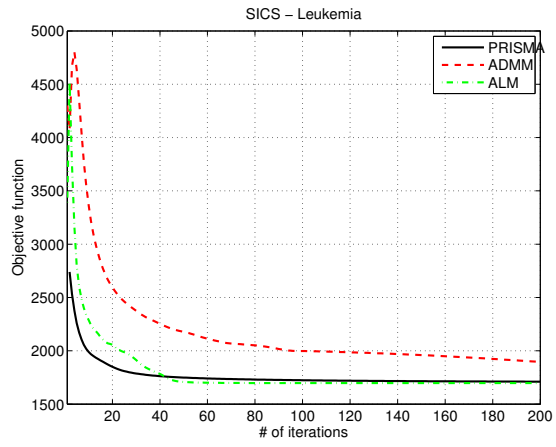


Figure 7: Comparison of the objective values as a function of the iterations.